The State of Cloud Storage

2013 Industry Report

A Benchmark Comparison
of Performance, Availability and Scalability

Executive Summary

In the last year, Cloud Storage Providers (CSPs) delivered over an exabyte of data under contract. This astonishing growth heralds a new era for how storage is delivered and consumed. The CSP market has seen new providers, new functionality, improved platform maturity, increased adoption in the enterprise and significant price reductions. Amazon's S3 service crossed the 1-trillion-object mark[1], a significant indication of platform maturity, and all major CSPs dropped their prices multiple times. As public cloud storage becomes critical to enterprise infrastructure, independent testing and monitoring of these systems is becoming all the more critical to IT.

Nasuni publishes this annual report to share the information that we assemble in order to properly evaluate CSPs for our own use. Cloud storage is a key component of our storage-as-a-service solution, much like the traditional system vendors leverage commodity hard drives in their storage arrays. Nasuni must monitor every change, improvement or update within the CSP market, and constantly evaluate which CSPs best enable us to provide the highest quality service to our customers.

In our 2012 report, tests demonstrated that Amazon S3 and Microsoft Azure Blob Storage were the two strongest players in the market and that, although other offerings showed potential, they had not yet matured enough for use in enterprise storage solutions.

This year, our tests revealed that Microsoft Azure Blob Storage has taken a significant step ahead of last year's leader, Amazon S3, to take the top spot. Across three primary tests (performance, scalability and stability), Microsoft emerged as a top performer in every category. Even though Azure has not been in the market as long as S3 and does not have nearly the same volume of storage under management, it is clear that Microsoft's investments in Azure are starting to pay dividends and that the technology they are providing to the market is second to none.

Overall, the CSPs tested this year demonstrated clear advancements over last year, including improved performance and fewer errors. It is clear that the minimum bar is moving upward, which is excellent news for the cloud storage market as a whole. As more CSPs mature into enterprise-class cloud storage providers, organizations and vendors will be able to leverage competitive advancements in price and technology to improve their overall storage infrastructure.

[1] **http://aws.typepad.com/aws/2012/06/amazon-s3-the-first-trillion-objects.html**

Why does Nasuni publish this report?

Nasuni provides storage infrastructure delivered to customers as a service and cloud storage is an essential component of the solution. We use cloud storage in much the same way that traditional storage vendors like EMC or NetApp use a hard drive - as a building block of a larger, more complete solution.

Customers use Nasuni to solve some of their most complex storage challenges, including dealing with explosive data growth, providing a uniform storage infrastructure across many offices, and making critical data available in every office and on mobile devices.

To solve these higher-level storage challenges, Nasuni works to identify and leverage components with capabilities far beyond what the traditional hard drive can provide. Cloud storage offers a unique and advanced set of benefits, including infinite scalability and global access. Still, it is also a relatively new technology in a market that is rapidly evolving. In order to take advantage of such technology and ensure quality, Nasuni engineers conduct frequent, ongoing testing and monitoring of Cloud Storage Providers.

This report represents a detailed annual review of Nasuni's latest findings. It is important to keep in mind that, while we do publish our findings, we conduct these tests primarily to identify CSPs that meet the needs of Nasuni and its customers. They do not necessarily represent the optimal metrics for any provider tested. However, these independent test results should be of interest to any organization that is considering the use of cloud storage as a component within a larger IT infrastructure.

Comparison Metrics

Nasuni partners with Cloud Storage Providers to achieve the best possible product at the most cost-effective price. Organizations considering cloud storage as part of their storage infrastructure should consider these same tradeoffs when comparing CSPs. Nasuni evaluates three key components of each CSP's offering:

- Functionality
- Price
- Performance

This annual report focuses primarily on reporting data from performance testing and touches only briefly on how Nasuni evaluates CSPs on functionality and price to provide context.

Functionality

While most interactions that an enterprise has with CSPs consist of simple API commands (GET, PUT and DELETE), an organization should consider a broader range of functionality when comparing cloud storage providers. Many companies today are global operations with offices around the world in a wide variety of localities, from major metropolitan areas to remote villages. To serve such users, cloud service providers need to maintain access points around the world and support meaningful cross-geography replication. Two copies of a file in a single datacenter is not geographic redundancy. In addition, organizations that expect to make meaningful use of cloud storage in their environment should also evaluate features of potential providers such as their API-based account creation and account management processes, availability of libraries and software to access data, the sophistication of their billing schemes and other aspects that help operations teams to ensure a smooth experience for their users and applications.

Price

Cloud storage architecture is fundamentally different from traditional storage; consequently, it is also priced differently from conventional storage. Instead of charging a price per raw TB (as with traditional storage hardware),most CSPs charge based on

GB stored per month. However, pricing is typically more complicated than a simple count of GB per month, often adding compute costs (to process API commands) and network costs (to move data to and from the cloud storage).

While this pricing model is cost effective because it charges customers only for the resources that they use, it makes predicting future costs a complex endeavor due to the variability of applications and use-cases. Although some vendors provide tools to help estimate costs, every customer's use-case is unique, so one-size-fits-all tools provide poor predictions. Unless the organization is working with a provider that offers a simplified pricing scheme, it is best to conduct initial tests with a minimal investment and then extrapolate from those results to develop a more accurate pricing estimation model.

Price itself is a very small part of a CSP comparison and may be the last part of a decision. Commodity offerings combined with competitive activity are driving costs down rapidly, however functionality and performance still vary significantly. When evaluating a CSP, remember, price is easy to change and negotiate – functionality and performance are not.

Performance

Performance is the primary yardstick by which Nasuni measures any publicly available CSP, testing the operation and stability of CSPs over long periods of time. In fact, Nasuni has been testing and comparing CSPs since 2009. Before considering any CSP for use in a production environment, it must meet minimum performance benchmarks across three areas:

Write/Read/Delete Benchmark: This simple test measures the raw ability of each CSP to handle thousands of writes, reads and deletes (W/R/D). We test each CSP with files of varying sizes:

- 1 KB
- 10 KB
- 100 KB
- 1 MB
- 10 MB
- 100 MB
- 1 GB

using varying levels of concurrency:

- 1 Thread
- 10 Threads
- 25 Threads
- 50 Threads

The Write/Read/Delete benchmark test runs for twelve hours, using multiple testing machine instances and several non-serial test runs to reduce the likelihood that external network issues could bias the results.

Availability: This test takes place over a 30-day period and measures each CSP's response time to a single W/R/D process at 60-second intervals:

- Write a randomly generated 1 KB file
- Read a randomly selected previously written file
- Delete a selected file

Reading and deleting a random file forces each CSP to prove their ability to be responsive to all of the data, all of the time, and not merely to the last piece of cached data. This test calculates the entire time required to complete the three requests, including any required retries. This ensures examination of not only responsiveness but also of CSP reliability and latency.

Scalability: Similar to the availability test, this is also an extended test that measures each CSP's ability to perform consistently as the number of objects under management increases. Performance under increasing object counts is often the Achilles heel of a cloud storage system, and this test measures each CSP's ability to maintain performance levels as the total number of objects stored in a single container increases to hundreds of millions.

Methodology

Due to dynamics in the marketplace, the list of platforms evaluated continues to change from year to year. The CSPs tested this year are Amazon S3, Microsoft Azure Blob Storage, Google Cloud Storage, HP Cloud Object Storage, and Rackspace Cloud Files. While many cloud storage platforms are publicly available, currently only these five platforms offer the combination of functionality, market experience and price that Nasuni requires for our solution. This includes three CSPs from last year (Amazon, Microsoft and Rackspace) and two newcomers (Google and HP).

Nasuni engineers conducted all tests between November 2012 and January 2013 using simple virtual machines across most of the major cloud-compute platforms. Each CSP was tested by using three "outside" machines (for example, Amazon EC2 does not test Amazon S3) spread throughout the eastern region of the United States. Although the use of "inside" machines would likely produce the best possible results for any CSP, we chose not to test such configurations, in order to present a scenario that accurately matches how Nasuni uses cloud storage - where it is accessed from outside of the cloud itself. We ran all tests using a variety of times, locations, virtual machines and dates to eliminate external network bias.

Virtual machine specifications:

- Machine 1
  - RAM: 15-16 GB
  - vCPUs: 4
  - Operating system: Ubuntu 12.04, 64 bit - Ubuntu12.04 LTS (GNU/Linux 3.2.0-24-virtual x86_64)

- Machine 2
  - RAM: 4 GB
  - vCPUs: 2
  - Operating system: Ubuntu 12.04, 64 bit - Ubuntu 12.04 LTS (GNU/Linux 3.2.0-24-virtual x86_64)

Note: although local storage was available, no local storage was used during testing. Furthermore, CPU load was monitored during the process to ensure it was never the bottleneck.

The tests are designed to evaluate the performance of CSPs under file-server data. The tests use the same distribution of file sizes used by actual Nasuni enterprise customers across thousands of installations over several years. Customers primarily use Nasuni to replace aging NAS technologies in datacenters and distributed offices – for this reason, the file size distribution matches that of a typical enterprise file server. In addition, this distribution of file sizes closely matches a well-documented breakdown from a study conducted by University of Wisconsin and Microsoft Research[2]. The exact distribution of file sizes used was:

1KB	10KB	100KB	1MB	10MB	100MB	1GB
16.8%	24.6%	26.2%	9.7%	22.2%	0.4%	0.1%

Table 1: File size distribution of Nasuni customers

This year, we invited the CSPs to participate in the evaluation process. The tested companies were allowed to review the preliminary results, discuss the findings and provide feedback to Nasuni engineering teams. We learned that in some cases, the CSPs had imposed limits on the machines, limits on container performance or had actually upgraded their networks during the testing period – all of which affected the results. For those CSPs that requested it, we re-ran the performance tests a second time. However, Nasuni did not allow any CSPs to make changes that would not be available at no additional cost to a standard customer.

Note: HP results may be limited due to rate limitations imposed on our testing account.

[2] **Agrawal, Bolosky, Douceur and Lorch. "A Five-Year Study of File-System Metadata" Microsoft Research. 1 October 2007. http://research.microsoft.com/pubs/72896/fast07-final.pdf**

One interesting lesson from the review process was that some CSPs required the creation of brand new accounts to take advantage of improved infrastructure[3]. While simple for testing, this is unacceptable in the real world and represents a clear deviation from the traditional Infrastructure-as-a-Service (IaaS) model. CSPs and any other as-a-service vendors should strive to provide their customers with the best possible customer experience, and all upgrades should occur behind the scenes without customer intervention. For organizations planning to use a CSP, it is best to communicate your use-case to the appropriate provider teams to ensure that you are taking advantage of their most up-to-date offerings. Furthermore, IT must stay informed about new offerings from existing suppliers so that they can take advantage of the latest advancements.

Results

Write/Read/Delete Benchmark

Benchmark results varied significantly across CSPs, illustrating that parity does not yet exist in this market. Of all the different tests run on the CSPs, this simple test highlighted the differences best. For each test (Write/Read/Delete), the test evaluated 23 combinations of file sizes and thread counts as shown in Table 2.

File Size & Threads Tested		1 KB	10 KB	100 KB	1 MB	10 MB	100 MB	1 GB
Number of Concurrent Threads	1	✓	✓	✓	✓	✓	✓	✓
	10	✓	✓	✓	✓	✓	✓	
	25	✓	✓	✓	✓	✓		
	50	✓	✓	✓	✓	✓		

Table 2: File size and thread count combinations tested

The results are averaged based on the weighting of customer file-server data (Table 1) and are then indexed to the performance of the top performer. The results, therefore, compare all the CSPs to the performance of the top performer across all file sizes and thread counts. This allows hundreds of individual tests to be evaluated using a single benchmark metric. Detailed raw results by CSP are included in the appendix.

The results show that raw write/read/delete performance varies significantly as object sizes and thread counts vary. Specifically, small object sizes and smaller thread counts highlight the transactional overhead of any platform. The effect of transactional overhead becomes most noticeable during writes, which contains three steps:

- Preparation
- Transmission
- Acknowledgement

For small files, Transmission is only a small portion of the total transaction, so any inefficiency in the performance of Preparation and Acknowledgement has greater impact. Those object stores that are built with efficient Preparation and Acknowledgement steps perform best when handling small files.

As file sizes or thread counts increase, the time associated with transmission increasingly dominates the overall time associated with the transaction. Inefficiencies in Preparation or Acknowledgement become less and less critical. Many CSPs overly focus their efforts to improve the efficiency of the Transmission stage of the transaction and thus perform better under the load of larger object sizes or thread counts. This may be fine for use-cases like media archives, but for file-server data which is often dominated by small files, performance on small files is critical.

[3] MSDN Blogs: Windows Azure. "Windows Azure's Flat Network Storage and 2012 Scalability Targets" 2 November 2012. http://blogs.msdn.com/b/windowsazure/archive/2012/11/02/windows-azure-s-flat-network-storage-and-2012-scalability-targets.aspx

Write Benchmark

Microsoft was the top write performer. Furthermore, Microsoft outperformed all other CSPs on 14 of the 23 individual combinations tested, making it far and away the optimal write target for file-based data. More so than in any other benchmark-based test, Microsoft shows how strong its updated technology is in this write test. Amazon and HP are the strongest second contenders trailing behind Microsoft (Figure 1).

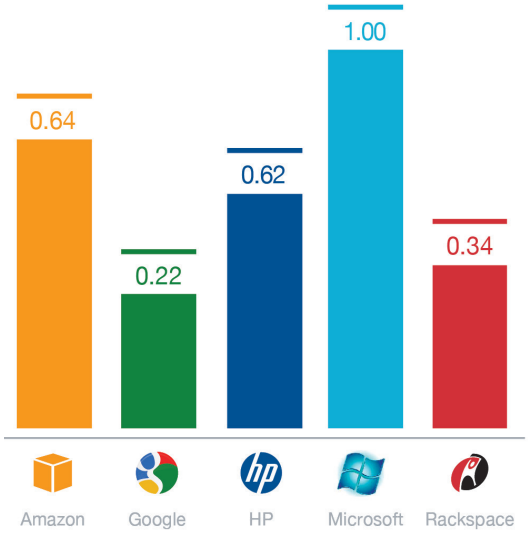

Figure 1: Indexed cloud storage write speed with all file sizes

The results for files larger than 1 MB show how much more closely all of the CSPs perform as transactional overhead becomes less significant than actual data transmission. In Figure 2, Amazon takes the top spot and the remaining providers show relative parity except for Google.

Read Benchmark

Read performance again shows Microsoft with a significant lead over its nearest competitor (Figure 3). However, Amazon no longer has the clear second

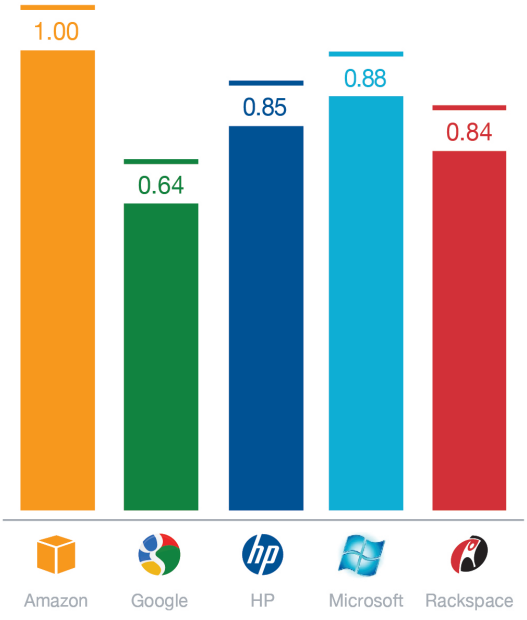

Figure 2: Indexed cloud storage write speed with files > 1MB

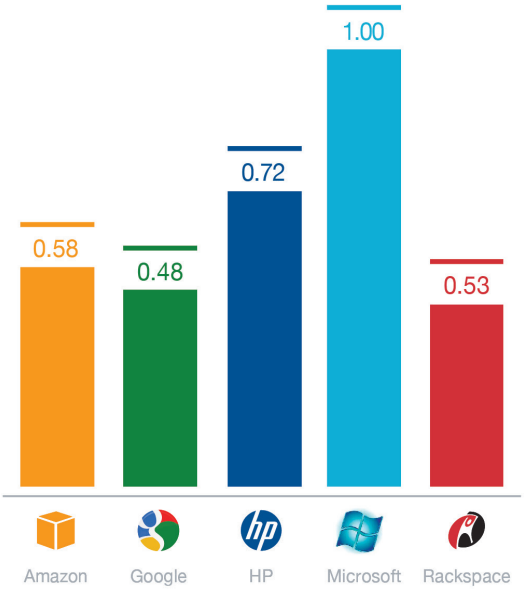

Figure 3: Indexed cloud storage read speed with all file sizes

www.nasuni.com
info@nasuni.com
United States: 1.800.208.3418
International: 1.508.433.6200

position from the write test. In fact, HP, a relative newcomer, edged out Amazon for the second spot. Even Google and Rackspace, which struggled by comparison on the write test, show much better relative benchmark performance when reading objects.

Unlike the results from the write benchmark, in Figure 4 Microsoft maintains its leadership position even as object sizes increase. While it still outperforms the other CSPs, the other platforms provide much closer performance than they do with smaller objects. Surprisingly, Amazon actually falls to fourth place behind both HP and Rackspace.

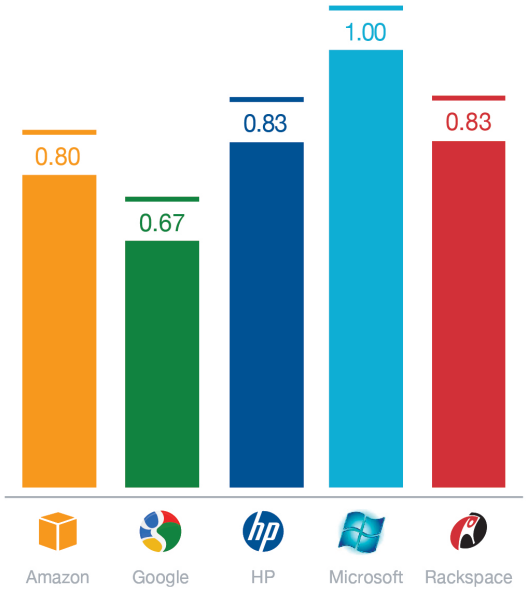

Figure 4: Indexed cloud storage read speed with files > 1MB

Delete Benchmark

Improving on the results from the other performance tests, Microsoft is more than twice as fast at deleting files as its nearest competitor. Amazon and HP share second position, well above both Google and Rackspace (Figure 5). Varying file sizes does not vary the results of this benchmark test.

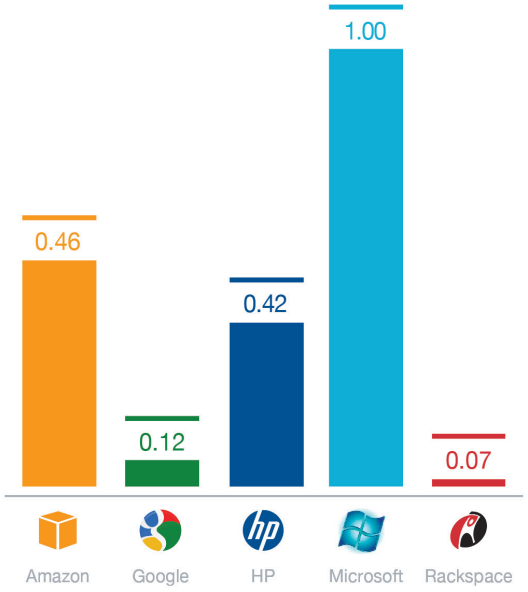

Figure 5: Indexed cloud storage delete speed

Availability

Availability is measured using the 'response time' metric that measures each CSP's response time to a single W/R/D process at 60-second intervals. Because 'response time' also includes any time associated with retries or delays, it is a more effective metric for availability than a simple ping test.

Microsoft performed the best in this test, averaging a response time of less than 0.5 seconds over a 30- day period. Amazon was the next closest, averaging just under 0.65 seconds, with third-place Rackspace averaging just under 1 second (Figure 6).

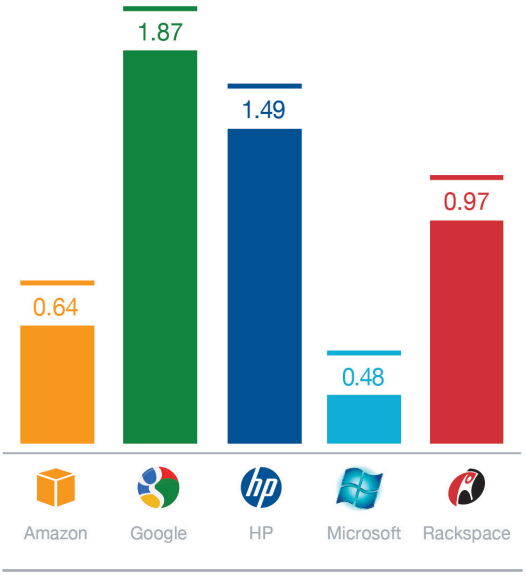

Figure 6: Average availability response time (seconds)

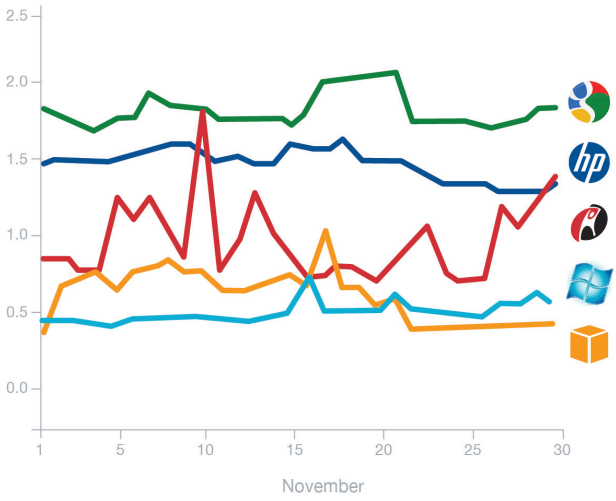

Figure 7: Average daily availability response time (seconds)

Examining the results over the month of testing also gives some insight into the variability of the numbers. While Microsoft and HP appear stable, Amazon and Google show a greater amount of variance in the results. Rackspace is the most inconsistent of all the CSPs, with meaningful swings visible day over day (Figure 7).

In addition to system and data availability, the test also measures overall uptime or percent of the time that the CSP is reachable. All players showed strong uptime percentages, with newcomer Google sharing the top spot with Amazon at 100 percent over the 30-day period (Figure 8).

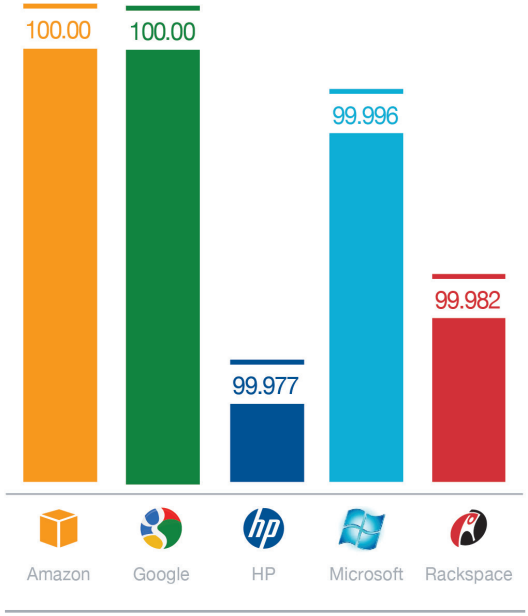

Figure 8: Average uptime

Scalability

As object counts increase, the performance of some CSPs degrade or becomes variable. Depending on CSP architecture, some systems are designed to scale across containers, not within them. This type of architectural limitation can become a significant bottleneck after months or even years of usage. An ideal scenario for

anyone seeking to leverage cloud storage is to partner with a CSP whose performance and responsiveness are unchanging regardless of the number of objects under management. Just as with traditional in-house storage, customers expect a consistent level of performance.

Under this test, all of the CSPs were loaded with new objects as quickly as possible up to 100 million objects or 30 days, whichever came first. The variance (Figure 9) represents how much the speed of loading objects changed over time, causing inconsistency and variability as objects were loaded.

Amazon, Microsoft and Google showed the lowest levels of variance, proving that no matter how many objects were loaded, performance did not depart significantly from their respective mean value. It is interesting to note that the two OpenStack platforms (HP and Rackspace) show the greatest variance as object counts increase. This may hint at architectural limitations, but without further testing and data, it is impossible to point to exactly what caused these elevated levels of variance and instability.

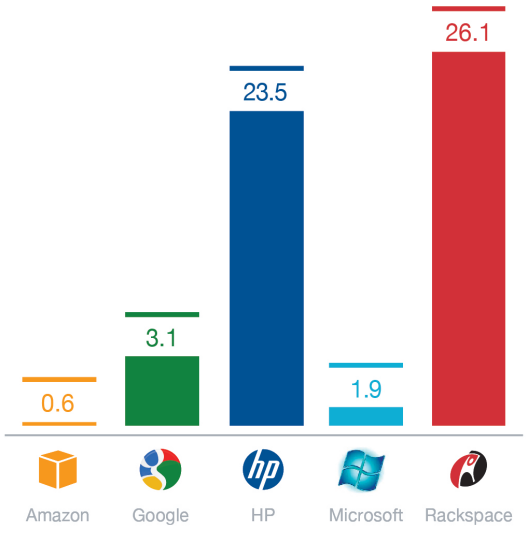

Figure 9: Variance during object scaling test (percent change in write speed)

Compared with last year's results, write and read error rates during scalability testing decreased significantly. While last year's results showed some error rates ranging from 1% - 60%, this year's numbers were all well below 1%.

During 100 million write attempts, Amazon, Microsoft and Google did not show a single write error. Rackspace had an error rate of 0.000001% and HP 0.000017% (Figure 10).

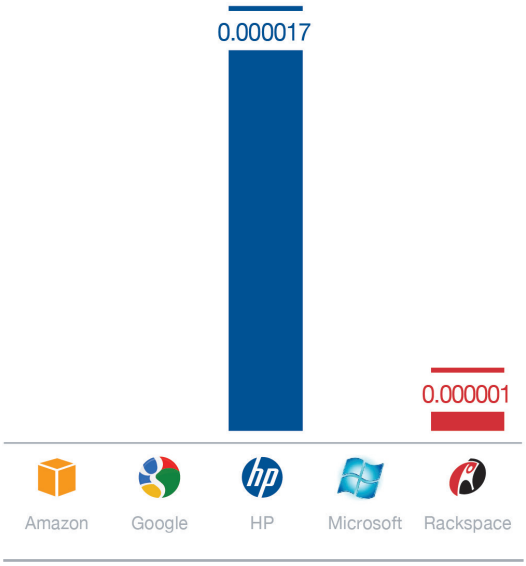

Figure 10: Percentage of write errors

During read attempts, only Microsoft resulted in no errors. Rackspace took the second spot with a read error rate of 0.0012% – significantly lower than last year's rate of 59%. HP had the highest read error rate, but was still below 0.01% (Figure 11).

www.nasuni.com
info@nasuni.com
United States: 1.800.208.3418
International: 1.508.433.6200

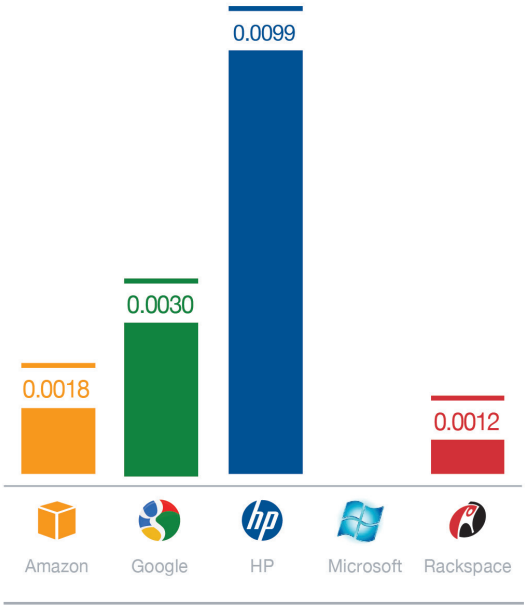

Figure 11: Percentage of read errors

Conclusion

Microsoft consistently performed better than the other CSPs in the tests, delivering the best Write/Read/Delete speeds across a variety of file sizes, the fastest response times and the fewest errors. Not only did Microsoft outperform the competition significantly during the raw performance tests, it was the only cloud storage platform to post zero errors during 100 million reads and writes. In those categories where Microsoft was not the top performer (uptime and scalability variance), it was a close second.

For these reasons, Microsoft has replaced Amazon to achieve the top performer position in the 2013 report. Cloud storage is a rapidly evolving market with new providers and new offerings entering all the time. The results in this report illustrate that product parity does not yet exist in this market. While offering "cloud storage" is relatively easy, delivering a high performing, reliable and scalable solution requires significant focus, advanced technology and continuous investment.

While Microsoft has secured the leadership position this year, it is quite possible that things could change again next year. In the two years this report has been published, there was a different leader each year. As presented in this paper, Amazon is still a strong player in this market, and is well positioned to continue to improve their platform producing better results. In addition, HP, a relative newcomer to the CSP market, showed strong performance in write and read tests, suggesting that improvements in stability could make them a key player. At Nasuni, we are committed to continuing to monitor, test and report on this dynamic market – watching closely to see who are the leaders and focusing on delivering the best complete cloud storage solution to our customers.

Disclaimer

The tests reported upon in this document are conducted by us using our own test tools under test conditions chosen by us. The test conditions were chosen by us to reasonably represent what our customers would experience using our Service with their representative environments and workloads. The tests have been designed by us to only look at the performance aspects of the CSPs that we believe are relevant to our customers – it is intentionally narrow in scope. Nasuni is not in the business of benchmarking CSPs, certifying test results or selling performance metrics. We have attempted to make sure the tests are fair and consistent within our selected parameters and have worked with several of the vendors to confirm our results. Our tests are not meant to indicate performance from each CSP under ideal conditions to the CSP, and, in any event, performance should only be one factor of many in a CSP selection process.

www.nasuni.com
info@nasuni.com
United States: 1.800.208.3418
International: 1.508.433.6200

Appendix – Detailed Benchmark Results by CSP

Amazon S3

Write Benchmark Results (MB/s)							
Threads	File Sizes						
	1 KB	10 KB	100 KB	1 MB	10 MB	100 MB	1 GB
1	0.02	0.13	1.10	5.82	12.29	13.80	13.05
10	0.15	1.34	10.48	52.00	58.01	66.80	
25	0.28	2.82	22.72	73.08	82.66		
50	0.27	2.84	25.57	80.94	89.07		

Read Benchmark Results (MB/s)							
Threads	File Sizes						
	1 KB	10 KB	100 KB	1 MB	10 MB	100 MB	1 GB
1	0.03	0.25	2.17	7.01	14.81	21.17	20.35
10	0.29	2.70	21.51	66.82	100.76	101.33	
25	0.76	6.89	45.46	83.72	102.11		
50	0.97	8.97	54.75	93.64	105.02		

Delete Benchmark Results (Obj/s)							
Threads	File Sizes						
	1 KB	10 KB	100 KB	1 MB	10 MB	100 MB	1 GB
1	30.33	33.32	34.70	35.43	34.64	36.82	14.16
10	350.31	331.70	378.30	375.02	248.54	195.02	
25	882.54	859.70	842.66	714.40	524.87		
50	1,000.86	991.70	1,041.73	976.58	411.72		

www.nasuni.com
info@nasuni.com
United States: 1.800.208.3418
International: 1.508.433.6200

Google Cloud Storage

Write Benchmark Results (MB/s)							
Threads	File Sizes						
	1 KB	10 KB	100 KB	1 MB	10 MB	100 MB	1 GB
1	0.00	0.01	0.07	0.68	4.42	12.53	16.22
10	0.01	0.08	0.68	6.58	41.11	59.30	
25	0.02	0.19	1.73	16.35	80.82		
50	0.04	0.37	3.48	31.73	107.36		

Read Benchmark Results (MB/s)							
Threads	File Sizes						
	1 KB	10 KB	100 KB	1 MB	10 MB	100 MB	1 GB
1	0.02	0.24	1.28	7.36	24.31	33.01	35.14
10	0.24	2.21	12.24	55.33	72.29	69.19	
25	0.56	5.21	34.47	69.17	69.11		
50	0.99	8.74	51.20	65.47	66.77		

Delete Benchmark Results (Obj/s)							
Threads	File Sizes						
	1 KB	10 KB	100 KB	1 MB	10 MB	100 MB	1 GB
1	7.25	7.66	7.65	7.83	7.34	7.42	5.57
10	79.88	79.29	78.79	70.94	66.08	49.50	
25	180.63	180.06	186.16	174.77	161.10		
50	337.62	338.64	352.24	334.03	246.26		

www.nasuni.com
info@nasuni.com
United States: 1.800.208.3418
International: 1.508.433.6200

HP Cloud Object Storage

Write Benchmark Results (MB/s)							
Threads	File Sizes						
	1 KB	10 KB	100 KB	1 MB	10 MB	100 MB	1 GB
1	0.02	0.16	2.85	10.72	13.41	13.69	13.95
10	0.14	1.95	17.55	47.69	47.13	49.24	
25	0.20	2.03	20.63	59.01	62.60		
50	0.22	2.26	21.44	58.75	49.41		

Read Benchmark Results (MB/s)							
Threads	File Sizes						
	1 KB	10 KB	100 KB	1 MB	10 MB	100 MB	1 GB
1	0.09	0.20	4.08	8.89	14.63	15.61	16.29
10	0.78	1.97	38.20	77.88	101.82	93.13	
25	1.75	4.92	64.58	97.99	103.15		
50	1.65	9.71	73.08	101.12	100.51		

Delete Benchmark Results (Obj/s)							
Threads	File Sizes						
	1 KB	10 KB	100 KB	1 MB	10 MB	100 MB	1 GB
1	48.67	57.40	48.13	52.97	42.29	20.28	5.65
10	358.17	393.95	383.88	402.42	391.55	115.81	
25	615.40	572.10	607.75	733.75	727.77		
50	760.65	752.52	743.79	913.45	674.31		

www.nasuni.com
info@nasuni.com
United States: 1.800.208.3418
International: 1.508.433.6200

Microsoft Azure Blob Storage

Write Benchmark Results (MB/s)							
Threads	File Sizes						
	1 KB	10 KB	100 KB	1 MB	10 MB	100 MB	1 GB
1	0.08	0.72	3.83	10.51	19.20		
10	0.38	3.60	23.62	53.94	42.31		
25	0.57	5.33	33.25	51.89	53.49		
50	0.67	6.40	25.60	56.34	59.17		

Read Benchmark Results (MB/s)							
Threads	File Sizes						
	1 KB	10 KB	100 KB	1 MB	10 MB	100 MB	1 GB
1	0.22	1.76	8.77	13.74	29.30		
10	1.61	11.08	44.92	74.93	102.97		
25	2.04	12.10	41.59	77.46	123.56		
50	1.89	12.81	42.76	92.58	145.73		

Delete Benchmark Results (Obj/s)							
Threads	File Sizes						
	1 KB	10 KB	100 KB	1 MB	10 MB	100 MB	1 GB
1	122.67	111.14	117.40	111.76	121.32		
10	880.85	872.91	882.55	841.09	607.20		
25	1,745.40	1,666.60	1,576.07	1,632.35	772.31		
50	1,671.51	1,713.77	1,567.42	1,306.09	867.50		

Note: Microsoft Azure Blog Storage does not support objects larger than 64MB in a single upload.

Rackspace Cloud Files

Write Benchmark Results (MB/s)							
Threads	File Sizes						
	1 KB	10 KB	100 KB	1 MB	10 MB	100 MB	1 GB
1	0.00	0.03	0.23	1.68	6.79	17.23	19.64
10	0.03	0.27	2.09	15.75	53.67	58.86	
25	0.07	0.64	5.20	37.38	98.63		
50	0.10	1.00	9.58	59.06	108.37		

Read Benchmark Results (MB/s)							
Threads	File Sizes						
	1 KB	10 KB	100 KB	1 MB	10 MB	100 MB	1 GB
1	0.02	0.19	1.44	8.09	21.29	34.09	35.99
10	0.19	1.87	13.87	57.64	80.75	74.99	
25	0.44	4.60	33.58	90.86	103.34		
50	0.87	8.58	63.63	100.50	119.24		

Delete Benchmark Results (Obj/s)							
Threads	File Sizes						
	1 KB	10 KB	100 KB	1 MB	10 MB	100 MB	1 GB
1	8.59	8.79	8.26	8.31	8.72	9.47	6.27
10	76.32	77.56	70.02	77.05	73.63	47.62	
25	96.61	99.21	98.00	98.34	86.10		
50	98.31	96.17	94.80	98.63	83.70		

www.nasuni.com
info@nasuni.com
United States: 1.800.208.3418
International: 1.508.433.6200

About Nasuni

Files are everywhere and they are a pain. Nasuni eliminates this pain forever by delivering file storage for distributed enterprises using a combination of cloud capacity, Nasuni software, and NAS appliances. Nasuni gives customers unlimited storage with built-in data protection and DR, secure global file sharing and mobile access, all managed from a single web console. Nasuni is cloud-based NAS for the distributed enterprise.

Our team is made up of enterprise storage, security and networking industry veterans with a shared vision of transforming the way enterprise organizations view data storage. We believe that storage should be as easy to purchase, consume and manage as the electricity that keeps the lights on.

Nasuni Corporation

www.nasuni.com
info@nasuni.com
United States: 1.800.208.3418
International: 1.508.433.6200
NAS_WP_10.14